

# It's LeVAsa not LevioSA!

## Latent Encodings for Valence-Arousal Structure Alignment

Surabhi S. Nath\*  
IIIT Delhi, India  
surabhi16271@iiitd.ac.in

Vishaal Udandaraao\*  
IIIT Delhi, India  
vishaal16119@iiitd.ac.in

Jainendra Shukla  
IIIT Delhi, India  
jainendra@iiitd.ac.in

### Abstract

In recent years, great strides have been made in the field of affective computing. Several models have been developed to represent and quantify emotions. Two popular ones include (i) categorical models which represent emotions as discrete labels, and (ii) dimensional models which represent emotions in a Valence-Arousal (VA) circumplex domain. However, there is no standard for annotation mapping between the two labelling methods. We build a novel algorithm for mapping categorical and dimensional model labels using annotation transfer across affective facial image datasets. Further, we utilize the transferred annotations to learn rich and interpretable data representations using a variational autoencoder (VAE). We present “LeVAsa”, a VAE model that learns implicit structure by aligning the latent space with the VA space. We evaluate the efficacy of LeVAsa by comparing performance with the Vanilla VAE using quantitative and qualitative analysis on two benchmark affective image datasets. Our results reveal that LeVAsa achieves high latent-circumplex alignment which leads to improved downstream categorical emotion prediction. The work also demonstrates the trade-off between degree of alignment and quality of reconstructions.

### Keywords

Affective computing, representation learning, variational autoencoder, valence, arousal, circumplex model

### ACM Reference Format:

Surabhi S. Nath, Vishaal Udandaraao, and Jainendra Shukla. 2021. It's LeVAsa not LevioSA! Latent Encodings for Valence-Arousal Structure Alignment. In *8th ACM IKDD CODS and 26th COMAD (CODS COMAD 2021)*, January 2–4, 2021, Bangalore, India. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3430984.3431037>

### 1 Introduction

Emotions are intrinsic characteristics of mammals, particularly overt in human behaviour [5, 17, 29]. Intelligent systems must employ means to incorporate emotions for a more natural interaction [31]. This surge for “emotional intelligence” has evolved into the field of affective computing, which by definition encompasses the creation of and interaction with machines that can sense, recognize,

respond to, and influence emotions [32]. Several models of emotion have been developed over the years, which are considered as the backbone of affective computing [11, 12, 25, 40]. Among these models, a popular choice is the Categorical Model which describes six basic discrete emotions, namely, happiness, anger, disgust, sadness, fear, and surprise [8]. However, this model failed to capture relations between the discrete emotions. Moreover, there is a lack of consistency in the choice of these fundamental emotions [7]. As a result, Russell & Mehrabian [35] developed the Dimensional Model which suggests that each emotional state can be defined in terms of Valence (pleasure of an emotion), Arousal (energy of an emotion) and Dominance (controlling nature of an emotion). The Dominance dimension is commonly ignored since the valence-arousal (VA) dimensional model was shown to possess adequate reliability, convergent validity, and discriminant validity [36]. This led to the conceptualization of the Circumplex Model to represent affective states as a circle in a 2D bipolar VA space [34]. The VA variables are typically considered independent [10].

**Motivation:** The existence of different models of emotions result in a range of possible annotation strategies for affective data [9, 21, 24, 27, 28]. This poses two challenges: (i) building deep models on affective data, and (ii) drawing collective insights from multiple datasets having potentially different formats of annotations [6].

In this paper, we present a novel algorithm for mapping annotations of the Categorical Model to those of the Dimensional Model through annotation transfer across affective facial image datasets. The subsequent step following annotation mapping is to obtain meaningful representations. With the increased use of deep neural networks and generative models, there have been significant advances in emotion modelling and affective computing [13, 19, 33]. Variational Autoencoders (VAEs) [20] are known to yield disentangled latent representations and generate new data samples [15, 16, 39]. They have been used extensively in affective computing to represent text, audio, image and electroencephalography (EEG) data [22, 41]. Applying VAEs on affective facial images to obtain disentangled image representations can (i) provide high quality feature representations for downstream tasks [2, 30], and (ii) serve applications like facial editing and data augmentation [23]. In our study, we obtain interpretable features by aligning the latent space of a VAE with the VA space. We show that this improves performance in downstream tasks of affect classification and regression, as demonstrated on two benchmark affective image datasets.

Our major contributions are (i) an annotation transfer algorithm for label transfer between Categorical and Dimensional models of emotion and (ii) a regularised VAE model “LeVAsa” (Latent Encodings for Valence-Arousal Structure Alignment) that yields an interpretable latent space with an implicit structure aligned with the VA space.

\*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
CODS COMAD 2021, January 2–4, 2021, Bangalore, India

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8817-7/20/01...\$15.00

<https://doi.org/10.1145/3430984.3431037>

## 2 Methods

Here we present our annotation transfer algorithm and VAE model architectures. Our code and models are publically available <sup>1</sup>.

### 2.1 Annotation Transfer Algorithm

For the task of annotation transfer between Categorical and Dimensional emotion models, we use an external reference dataset ( $D_r$ ) containing both discrete categorical emotion labels ( $l_i \in \{e_1, e_2, \dots, e_n\}$ , where  $e_1, e_2, \dots, e_n$  are the  $n$  discrete emotional labels) and valence, arousal values ( $v_i \in [l_{lim_v}, ulim_v]$ ,  $a_i \in [l_{lim_a}, ulim_a]$ , where  $l_{lim_v}, ulim_v, l_{lim_a}, ulim_a$  are the lower and upper limits for valence and arousal values respectively). Each data sample  $x_i \in D_r$  thus has an emotion label  $l_i$ , a valence value  $v_i$  and an arousal value  $a_i$ .  $D_r$  serves as the standard based on which continuous or discrete VA values can be sampled for data points in a working dataset ( $D$ ) with only emotion labels (Algorithm 1, Line 4), or conversely, the most likely emotion labels can be obtained for data points in dataset  $D'$  with only VA tuples (Algorithm 1, Line 6). The ellipse sizes (on average 3% of total area) ensure that the sampled VA values allow both sufficient variability and consistency within emotion classes.

---

#### Algorithm 1: Annotation transfer algorithm

---

**Input** : reference dataset  $D_r$ , discrete categorical emotion labels  $[l_1, l_2, \dots, l_{|D_r|}]$ , VA values  $[v_1, v_2, \dots, v_{|D_r|}]$  and  $[a_1, a_2, \dots, a_{|D_r|}]$ , working dataset  $D$  with discrete emotion labels, working dataset  $D'$  with VA tuples

**Output**: VA values for the working dataset  $D$ , discrete emotion labels for the working dataset  $D'$

- 1 Partition each sample  $x_i \in D_r$  into  $n$  groups ( $g_1, g_2, \dots, g_n$ ) based on discrete emotion labels  $l_i \in \{e_1, e_2, \dots, e_n\}$
  - 2 For each group  $g_j$ ,  $j \in \{1, 2, \dots, n\}$ , obtain the mean valence  $\mu_{vj}$ , standard deviation valence  $\sigma_{vj}$ , mean arousal  $\mu_{aj}$ , standard deviation arousal  $\sigma_{aj}$  values
  - 3 Generate ellipses  $c_j$ ,  $j \in \{1, 2, \dots, n\}$  for group  $g_j$  representing emotion  $e_j$  with centre  $(\mu_{vj}, \mu_{aj})$ , semi major axis  $\sigma_{vj}$  and semi minor axis  $\sigma_{aj}$
  - 4 To obtain VA values for data point  $x_k$  in  $D$  with label  $l_k$ , sample  $(v_k, a_k)$  from ellipse  $c_k$  as:  $v_k = x\sigma_{vk} + \mu_{vk}$ ,  $a_k = y\sigma_{ak} + \mu_{ak}$ , where  $x = \sqrt{r} \cos \theta$ ,  $y = \sqrt{r} \sin \theta$  and  $r \sim [0, 1]$ ,  $\theta \sim [0, 2\pi]$
  - 5 Convert  $(v_k, a_k)$  to discrete values by scaling and rounding-off if desired
  - 6 To obtain emotion label for sample  $x_k$  in  $D'$  with VA  $(v_k, a_k)$ , find ellipse  $c_k$  with centroid at least Euclidean distance from  $(v_k, a_k)$ , and assign  $e_k$  as most likely emotion
- 

### 2.2 VAE model architectures

We train a generative model with an interpretable latent space with an implicit structure given a raw distribution of affective face images. We employ variational autoencoder based models because of their simple training protocols and structured inductive priors. We compare two VAE models, Vanilla VAE and LeVAsa. This comparison is justified because: (1) No baseline VA structured latent

<sup>1</sup><https://github.com/vishaal27/LeVAsa>




Figure 1: Model Architecture

model exists, and (2) the Vanilla VAE features being unbiased (owing to no explicit supervision) may perform better on the downstream tasks. The latent space for both models was constructed to comprise of three chunks. Figure 1 depicts our model architectures.

For the Vanilla VAE, no explicit alignment was imposed on the latent space, whereas for LeVAsa, we take inspiration from recent work [4, 18] and model the latent space as follows:

- $Z_v$  – subspace consisting of valence attributes  $z_v$  that learn to encode the valence features of image samples
- $Z_a$  – subspace consisting of arousal attributes  $z_a$  that learn to encode the arousal features of image samples
- $Z_z$  – subspace consisting of other miscellaneous generative attributes  $z_z$  that are required for high-fidelity reconstruction of the input data distribution.

Given a dataset of  $N$  affective images  $X = \{x_1, x_2, \dots, x_N\}$ , our VAE backbone consists of encoder  $f_\theta$  and decoder  $g_\phi$  given by:

$$\begin{aligned} z_{v_i}, z_{a_i}, z_{z_i} &= f_\theta(x_i), x_i \in X \\ \hat{x}_i &= g_\phi(z_{v_i}, z_{a_i}, z_{z_i}) \\ z_{v_i} &\in Z_v, z_{a_i} \in Z_a, z_{z_i} \in Z_z \end{aligned}$$

We train the Vanilla VAE with a simple reconstruction loss along with a modified Kullback-Leibler (KL) loss (Eq. 1). We induce a  $\mathcal{N}(\mathbf{0}, \mathbf{I})$  prior on all three attributes  $z_v, z_a$  and  $z_z$ .

$$\mathcal{L}_{KL} = \sum_{z \in \{z_v, z_a, z_z\}} KL(f_\theta(z|x) || \mathcal{N}(\mathbf{0}, \mathbf{I}))^2 \quad (1)$$

We employ the same backbone Vanilla VAE architecture for the LeVAsa model with two major modifications:

- (1) *Projection Heads*: We use two non-linear projection heads  $h_v$  and  $h_a$  which map the encoded valence and arousal representations  $z_v$  and  $z_a$  to the valence and arousal label space (giving label representations  $r_v$  and  $r_a$ ). The projections obtained are represented as follows:
$$r_{v_i} = h_v(z_{v_i}), r_{a_i} = h_a(z_{a_i})$$
- (2) *VA-regularization loss*: To impose an explicit alignment of the  $z_v$  and  $z_a$  attributes with the VA ground truth factors,

we introduce a VA-regularization loss as follows:

$$\mathcal{L}_C = \sum_{i=0}^N \left( \mathcal{L}(r_{v_i}, v_i) + \mathcal{L}(r_{a_i}, a_i) \right) \quad (2)$$

where  $\mathcal{L}$  takes the form of MSE for continuous and BCE for discrete annotation types. MSE/BCE are design choices and can be replaced by suitable likelihood-based loss functions.

The overall optimization objective for the LeVAsa model is:

$$\mathcal{L}_{total} = \mathcal{L}_R + \lambda_{KL} \mathcal{L}_{KL} + \lambda_C \mathcal{L}_C \quad (3)$$

where  $\lambda_{KL}$  and  $\lambda_C$  are hyperparameters.

### 2.3 Datasets

We use the following datasets in our experiments.

*Annotation Transfer: AffectNet*

- **AffectNet** [26] is the largest facial expression dataset, with over 420,000 annotated images and contains both continuous VA annotations and discrete emotional labels. The dataset also incorporates a wide diversity in gender, age and ethnicity, hence is an ideal choice for the reference dataset in the annotation transfer algorithm (Algorithm 1). The generated ellipses are shown in Figure 2.

*Model Training: IMFDB, AFEW-VA*

- **IMFDB** [38] contains around 34,000 annotated zoomed-in facial images of 100 Indian actors, with only emotional labels and no VA supervision. Continuous and discrete VA supervision for IMFDB is obtained from annotation transfer using AffectNet. This is particularly well suited due to the similar nature of images in IMFDB and AffectNet datasets.
- **AFEW-VA** [21] on the other hand, contains around 24,000 annotated images from videos of real world scenes of approximately 600 actors with only discrete VA values and no discrete emotional labels.

The different nature of IMFDB and AFEW-VA datasets allow us to analyse and compare model performance based on different factors including image type (zoomed in faces/video scenes) and annotation type (discrete VA supervision/continuous VA supervision).




Figure 2: Ellipses from AffectNet for annotation transfer

## 3 Experiments

We evaluate model performance through a series of qualitative and quantitative experiments. This enables comparisons based on three aspects: (i) architecture (Vanilla VAE vs LeVAsa), (ii) dataset (IMFDB vs AFEW-VA), and (iii) nature of annotations (Continuous VA vs Discrete VA). Altogether, we train five models: (i) Vanilla VAE on

IMFDB, (ii) LeVAsa on IMFDB with continuous VA annotations, (iii) LeVAsa on IMFDB with discrete VA annotations, (iv) Vanilla VAE on AFEW-VA, (v) LeVAsa on AFEW-VA with discrete VA annotations.

### 3.1 Latent-Circumplex Alignment

We measure the alignment of LeVAsa's  $Z_v \cup Z_a$  latent space with the VA ground truths using normalized Euclidean and Manhattan distance metrics for continuous annotations, and Cross Entropy measure for discrete annotations. This helps quantify the degree of latent-circumplex alignment. For the Vanilla VAE, we determine the  $z_v$  and  $z_a$  chunks heuristically by considering the two latent chunks which aligned best with the corresponding valence and arousal ground truths. Further, we reduce the dimensionality of the  $z_v$  and  $z_a$  latent chunks and plot them alongside the ground truth to replicate the circumplex representation.

It is found that LeVAsa outperformed Vanilla VAE for both continuous and discrete annotations (Table 1). This clearly exhibits the superior latent-circumplex alignment achieved by LeVAsa. For discrete annotations, in case of AFEW-VA, the difference between the cross entropy measures of the Vanilla VAE and LeVAsa is greater than in case of IMFDB. This could be attributed to the different image types in both datasets. The circumplex plots (Figure 3) for LeVAsa reveal reduced variance and increased alignment with true labels. This validates the quantitative results in Table 1.


Table 1: Alignment

IMFDB	Valence		Arousal		Combined	
	MSE	MAE	MSE	MAE	MSE	MAE
Vanilla VAE	1.83	0.29	1.49	0.26	3.31	0.55
LeVAsa	0.14	0.14	0.06	0.09	0.2	0.23

(a) Continuous

Model	IMFDB	AFEW-VA
Vanilla VAE	8.9	8.9
LeVAsa	6.63	2.54

(b) Discrete



(a) IMFDB Cont. (b) IMFDB Disc. (c) AFEW-VA

Figure 3: Circumplex Representation

To gain further insights, we assess the regressive power of  $z_v$  and  $z_a$  by their ability to predict the corresponding VA ground truths. We use Multi Layer Perceptron (MLP) Regression. This analysis applies to continuous annotations hence was conducted only on the models trained on IMFDB dataset with continuous VA values.

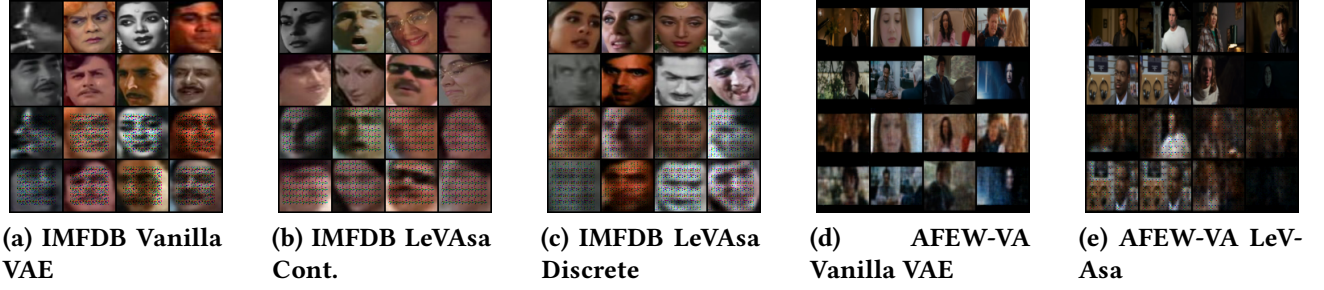


Figure 4: VAE Reconstruction from the five models

Table 2: VA Regressive Power

Axis	Model	MSE	MAE	EV	$R^2$
Valence	Vanilla VAE	0.256	0.420	-0.011	-0.012
	LeVAsa	0.251	0.414	0.016	0.015
Arousal	Vanilla VAE	0.092	0.242	-0.022	-0.048
	LeVAsa	0.074	0.224	0.095	0.086

It is observed that the MSE and MAE values computed for LeVAsa were lower by 2.25% and 1.42% as compared to Vanilla VAE for valence, and lower by 19.13% and 7.18% as compared to Vanilla VAE for arousal (Table 2). Furthermore, the goodness of fit metrics (explained variance and  $R^2$ ) showed better performance in the case of LeVAsa. These results further strengthen our hypothesis.

### 3.2 Categorical Emotion Predictive Power

We predict the discrete emotion labels using different combinations of latent representations obtained from Vanilla VAE and LeVAsa (Table 3). Due to lack of discrete emotion labels in the AFEW-VA dataset, it was excluded from this analysis. We randomized the data splits across Continuous and Discrete experiments to ensure an unbiased setup. Model performance is evaluated using classification accuracy. We utilize a simple one-layered MLP to ensure that the accuracy is a direct measure of representation quality and not influenced by the complexity of the classifier.

It is seen that LeVAsa has significantly better predictive power as compared to the Vanilla VAE. Moreover, for LeVAsa, the VA chunks alone are more informative in emotion prediction as compared to  $z_v \oplus z_a \oplus z_z$  chunks altogether. Also, the improvement in classification accuracy by employing LeVAsa in place of Vanilla VAE can be compared under the continuous and discrete settings. This reveals that LeVAsa representations from the model trained with discrete annotations and BCE loss (Eq. 2) proves to be better at classifying emotion labels. This is due to the discrete nature of emotion labels which correlate well with the model representations.

### 3.3 Reconstruction Quality

VAE models are prone to posterior collapse, producing unreliable reconstructions [14, 37]. Thus, along with analyses of the latent representations, we study the quality of reconstructions (Figure 4).

It is observed that the quality of the reconstructed faces is slightly compromised in the case of LeVAsa as compared to Vanilla VAE. This can be attributed to the slightly higher variance of the learnt

Table 3: Categorical emotion predictive power using vanilla VAE and LeVAsa models. All reported scores are accuracies ( $\oplus$  represents vector concatenation)

Annotation Type	Chunk Combination	Vanilla VAE (V)	LeVAsa (L)	Difference= L - V (in %)
Continuous	$z_v$	0.29	0.36	7
	$z_a$	0.32	0.35	3
	$z_z$	0.32	0.36	4
	$z_v \oplus z_a$	0.32	0.38	6
	$z_v \oplus z_a \oplus z_z$	0.29	0.33	4
Discrete	$z_v$	0.30	0.35	5
	$z_a$	0.27	0.30	3
	$z_z$	0.24	0.30	6
	$z_v \oplus z_a$	0.25	0.33	8
	$z_v \oplus z_a \oplus z_z$	0.26	0.30	4

LeVAsa decoding distribution [1, 15]. By Shannon’s rate-distortion theory [3], there is a trade-off between the distortion (reconstruction quality) and rate (representation quality). Since we are imposing an explicit compression bottleneck on the latent space, it is expected that the reconstruction quality is slightly compromised in order to achieve better interpretability of latent representations.

## 4 Conclusion

In this paper, we have developed an annotation-transfer algorithm for mapping between Categorical and Dimensional emotion model annotations. Using them, we generated interpretable image features with a VA-regularized VAE model called LeVAsa. We conducted a series of evaluation tasks to verify and validate our experiments and compare performance based on three factors: (i) architecture (Vanilla VAE vs LeVAsa), (ii) dataset (IMFDB vs AFEW-VA), and (iii) nature of annotations (Continuous VA vs Discrete VA). The results showed that the LeVAsa model obtains robust and interpretable representations enabling improved downstream affective task performance. In the future, we hope to extend the annotation-transfer algorithm to action-unit annotations, and test the expressiveness of the representations by performing latent traversals for data augmentation and facial editing.

## Acknowledgments

This work was supported by the Infosys Center for Artificial Intelligence at IIT Delhi, India.

## References

- [1] Alexander Alemi, Ben Poole, Ian Fischer, Joshua Dillon, Rif A Saurous, and Kevin Murphy. 2018. Fixing a broken ELBO. In *International Conference on Machine Learning*. 159–168.
- [2] Yoshua Bengio, Aaron Courville, and Pascal Vincent. 2013. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence* 35, 8 (2013), 1798–1828.
- [3] Toby Berger. 2003. Rate-distortion theory. *Wiley Encyclopedia of Telecommunications* (2003).
- [4] Sarthak Bhagat, Vishaal Udandara, and Shagun Uppal. 2020. DisCont: Self-Supervised Visual Attribute Disentanglement using Context Vectors. *arXiv preprint arXiv:2006.05895* (2020).
- [5] Charles Darwin and Phillip Prodger. 1998. *The expression of the emotions in man and animals*. Oxford University Press.
- [6] Luna De Bruyne, Pepa Atanasova, and Isabelle Augenstein. 2019. Joint Emotion Label Space Modelling for Affect Lexica. *arXiv preprint arXiv:1911.08782* (2019).
- [7] Paul Ekman and Daniel Cordaro. 2011. What is meant by calling emotions basic. *Emotion review* 3, 4 (2011), 364–370.
- [8] Paul Ekman and Wallace V Friesen. 1971. Constants across cultures in the face and emotion. *Journal of personality and social psychology* 17, 2 (1971), 124–129.
- [9] C Fabian Benitez-Quiroz, Ramprakash Srinivasan, and Aleix M Martinez. 2016. Emotionet: An accurate, real-time algorithm for the automatic annotation of a million facial expressions in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5562–5570.
- [10] Lisa Feldman Barrett and James A Russell. 1998. Independence and bipolarity in the structure of current affect. *Journal of personality and social psychology* 74, 4 (1998), 967–984.
- [11] Jonathan Gratch, Stacy Marsella, Ning Wang, and Brooke Stankovic. 2009. Assessing the validity of appraisal-based models of emotion. In *2009 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*. IEEE, 1–8.
- [12] Stephan Hamann. 2012. Mapping discrete and dimensional emotions onto the brain: controversies and consensus. *Trends in cognitive sciences* 16, 9 (2012), 458–466.
- [13] Jing Han, Zixing Zhang, and Bjorn Schuller. 2019. Adversarial training in affective computing and sentiment analysis: Recent advances and perspectives. *IEEE Computational Intelligence Magazine* 14, 2 (2019), 68–81.
- [14] Junxian He, Daniel Spokoyny, Graham Neubig, and Taylor Berg-Kirkpatrick. 2019. Lagging inference networks and posterior collapse in variational autoencoders. *arXiv preprint arXiv:1901.05534* (2019).
- [15] Irina Higgins, Loïc Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew M Botvinick, Shakir Mohamed, and Alexander Lerchner. 2017. beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In *ICLR*. 1–22.
- [16] Qiyang Hu, Attila Szabó, Tiziano Portenier, Paolo Favaro, and Matthias Zwicker. 2018. Disentangling factors of variation by mixing them. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3399–3407.
- [17] Carroll E Izard. 2013. *Human emotions*. Springer Science & Business Media.
- [18] Ananya Harsh Jha, Saket Anand, Maneesh Singh, and VSR Veeravasarpur. 2018. Disentangling factors of variation with cycle-consistent variational autoencoders. In *European Conference on Computer Vision*. Springer, 829–845.
- [19] Baani Leen Kaur Jolly, Palash Aggrawal, Surabhi S Nath, Viresh Gupta, Manraj Singh Grover, and Rajiv Ratn Shah. 2019. Universal EEG Encoder for Learning Diverse Intelligent Tasks. In *2019 IEEE Fifth International Conference on Multimedia Big Data (BigMM)*. IEEE, 213–218.
- [20] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
- [21] Jean Kossaihi, Georgios Tzimiropoulos, Sinisa Todorovic, and Maja Pantic. 2017. AFEW-VA database for valence and arousal estimation in-the-wild. *Image and Vision Computing* 65 (2017), 23–36. Elsevier.
- [22] Siddique Latif, Rajib Rana, Junaid Qadir, and Julien Epps. 2017. Variational autoencoders for learning latent representations of speech emotion: A preliminary study. *arXiv preprint arXiv:1712.08708* (2017).
- [23] Alexandra Lindt, Pablo Barros, Henrique Siqueira, and Stefan Wermter. 2019. Facial expression editing with continuous emotion labels. In *2019 14th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2019)*. IEEE, 1–8.
- [24] Patrick Lucey, Jeffrey F Cohn, Takeo Kanade, Jason Saragih, Zara Ambadar, and Iain Matthews. 2010. The extended cohn-kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. In *2010 IEEE computer society conference on computer vision and pattern recognition-workshops*. IEEE, 94–101.
- [25] Stacy Marsella, Jonathan Gratch, Paolo Petta, et al. 2010. Computational models of emotion. *A Blueprint for Affective Computing-A sourcebook and manual* 11, 1 (2010), 21–46.
- [26] Ali Mollahosseini, Behzad Hasani, and Mohammad H Mahoor. 2017. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing* 10, 1 (2017), 18–31.
- [27] Mihalís A Nicolaou, Hatice Gunes, and Maja Pantic. 2010. Audio-visual classification and fusion of spontaneous affective data in likelihood space. In *2010 20th International Conference on Pattern Recognition*. IEEE, 3695–3699.
- [28] Mihalís A Nicolaou, Hatice Gunes, and Maja Pantic. 2011. Continuous prediction of spontaneous affect from multiple cues and modalities in valence-arousal space. *IEEE Transactions on Affective Computing* 2, 2 (2011), 92–105.
- [29] Jaak Panksepp. 2004. *Affective neuroscience: The foundations of human and animal emotions*. Oxford University Press.
- [30] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. 2017. *Elements of causal inference*. The MIT Press.
- [31] Rosalind W Picard. 2000. *Affective computing*. The MIT Press.
- [32] Rosalind W Picard and Jonathan Klein. 2002. Computers that recognise and respond to user emotion: theoretical and practical implications. *Interacting with computers* 14, 2 (2002), 141–169.
- [33] Philipp V Rouast, Marc Adam, and Raymond Chiong. 2019. Deep learning for human affect recognition: insights and new developments. *IEEE Transactions on Affective Computing* (2019), 1–20.
- [34] James A Russell. 1980. A circumplex model of affect. *Journal of personality and social psychology* 39, 6 (1980), 1161–1178.
- [35] James A Russell and Albert Mehrabian. 1977. Evidence for a three-factor theory of emotions. *Journal of research in Personality* 11, 3 (1977), 273–294.
- [36] James A Russell, Anna Weiss, and Gerald A Mendelsohn. 1989. Affect grid: a single-item scale of pleasure and arousal. *Journal of personality and social psychology* 57, 3 (1989), 493–502.
- [37] Oleh Rybkin, Kostas Daniilidis, and Sergey Levine. 2020. Simple and Effective VAE Training with Calibrated Decoders. *arXiv preprint arXiv:2006.13202* (2020).
- [38] Shankar Setty, Moula Husain, Parisa Beham, Jyothi Gudavalli, Menaka Kandasamy, Radhesyam Vaddi, Vidyagouri Hemadri, JC Karure, Raja Raju, B Rajan, et al. 2013. Indian movie face database: a benchmark for face recognition under wide variations. In *2013 fourth national conference on computer vision, pattern recognition, image processing and graphics (NCVPRIPG)*. IEEE, 1–5.
- [39] Ankita Shukla, Sarthak Bhagat, Shagun Uppal, Saket Anand, and Pavan K. Turaga. 2019. Product of Orthogonal Spheres Parameterization for Disentangled Representation Learning. In *BMVC*. 1–13.
- [40] Jessica L Tracy and Daniel Randles. 2011. Four models of basic emotions: a review of Ekman and Cordaro, Izard, Levenson, and Panksepp and Watt. *Emotion Review* 3, 4 (2011), 397–405.
- [41] Chuhan Wu, Fangzhao Wu, Sixing Wu, Zhigang Yuan, Junxin Liu, and Yongfeng Huang. 2019. Semi-supervised dimensional sentiment analysis with variational autoencoder. *Knowledge-Based Systems* 165 (2019), 30–39.